

Formal Analysis of Vulnerabilities of Web Applications Based on SQL Injection^{*}

(Extended Version)

Federico De Meo¹, Marco Rocchetto², and Luca Viganò³

¹ Dipartimento di Informatica, Università degli Studi di Verona, Italy

² iTrust, Singapore University of Technology and Design, Singapore

³ Department of Informatics, King's College London, UK

Abstract. We present a formal approach for the analysis of attacks that exploit SQLi to violate security properties of web applications. We give a formal representation of web applications and databases, and show that our formalization effectively exploits SQLi attacks. We implemented our approach in a prototype tool called SQLfast and we show its efficiency on four real-world case studies, including the discovery of an attack on Joomla! that no other tool can find.

1 Introduction

Motivations. According to OWASP (the Open Web Applications Security Project [28]), *SQL injection (SQLi)* is the most critical threat for the security of web applications (*web apps*, for short), and MITRE lists improper SQLi neutralization as the most dangerous programming error [7]. SQLi was first defined in [14] but, also due to the increasing complexity of web apps, SQLis can still be very difficult to detect, especially by manual *penetration testing (pentesting)*.

A number of SQLi scanners have thus been developed to search for injection points and payloads, most notably *sqlmap* [34], which allows human pentesters to find SQLi vulnerabilities by testing the web app with different payloads, and *sqlninja* [35], which focuses on SQL server databases. The combination of the two provides the pentester with a powerful tool suite for SQLi detection. However, neither sqlmap nor sqlninja (nor other state-of-the-art vulnerability scanners) are able to detect vulnerabilities linked to logical flaws of web apps [12]. This means that even if a scanner can concretely discover a SQLi, it can't link SQLi to logical flaws that lead to the violation of a generic security property, e.g., the secrecy of data accessible only bypassing an authentication phase via a SQLi.

Moreover, determining that a web app is vulnerable to SQLi (and which payload to exploit) might not be enough for the app's overall security. Consider, for instance, a web app that relies on legacy code (when an update is not feasible, e.g., because the legacy code is a core part of the system). If a SQLi is found, an investigation should be performed to understand when the SQLi can be exploited

^{*} This work was carried out while Marco Rocchetto was at the Università di Verona.

and whether this compromises security. This investigation is carried out manually by the pentester in charge of identifying attack scenarios, thus potentially leading to additional omissions, errors and oversights in the security analysis.

A number of formal approaches for the security analysis of web apps, based on the *Dolev-Yao (DY) intruder model* [11], have been implemented recently, e.g., [1,3,5,32,38]. However, the DY model is typically used to reason about security protocols and the cryptographic operators they employ (e.g., for asymmetric or symmetric cryptography, modular exponentiation or exclusive-or) but abstracting away the contents of the payloads of the messages. As a consequence, these approaches cannot properly identify or exploit new SQLi payloads since reasoning about the contents of the messages is crucial to that end.

Contributions. In this paper, we present a formal approach for the analysis of attacks that exploit SQLi to violate security properties of web apps. We define how to formally represent web apps that interact with a database and how the DY intruder model can be extended to deal with SQLi.

In order to show that our formalization can effectively be used to detect security vulnerabilities linked to SQLi attacks, we have developed a prototype tool called *SQLfast (SQL Formal Analysis Tool)* and we show its efficiency by discussing four real-world case studies. Most notably, we use SQLfast to detect an attack on Joomla! which, to the best of our knowledge, no state-of-the-art SQLi scanner (e.g., sqlmap or sqlninja) can detect since they do not automatically link different attacks in one attack trace (i.e., they do not find logical flaws linked to SQLi attacks). Another key novel aspect of SQLfast is that it can detect complex attacks in which a first SQLi attack provides data for a second subsequent attack. We show that SQLfast allows us to exploit SQLi combining it with logical flaws of web apps to report sophisticated attack traces in a few seconds and can also deal with Second-Order SQLis, which are notoriously difficult to spot.

Note that we do not search for new SQLi payloads but rather we exploit attacks related to SQLi. This allows us to analyze how an intruder can violate a security property by exploiting one or more attacks related to a SQLi, e.g., credential bypass. Nevertheless, we also (automatically) test our attacks against the web app under analysis and then we use state-of-the-art tools (i.e., sqlmap and curl) to detect the actual payload of all the SQLis exploited.

Organization. In § 2, we discuss a concrete example that shows why we can't stop at the identification of a SQLi. In § 3, we give a categorization of SQLi vulnerabilities to highlight the main aspects that characterize SQLi, based on which, in § 4, we provide our formalization. In § 5, we discuss SQLfast and its application real-world case studies. We discuss related work in § 6 and conclude in § 7. The appendix contains a proof that the formalization of the database correctly handles all the SQLis categorized in § 3, as well as full details on our specifications and case studies.

2 Why can't we stop at the identification of a SQLi? The case of Joomla!

The identification of a SQLi entry point is generally considered as a satisfactory finish line when dealing with SQLi in web apps. So, one could ask: why not simply stop there (and why bother reading the rest of this paper)? The answer is that a SQLi can be a serious threat *only* if it can be exploited and *only* if it can be used for carrying out an attack. A full understanding of how a potential SQLi vulnerability can afflict the security of a web app is essential in order to implement proper countermeasures. For instance, consider Joomla! [23], a PHP-based Content Management System that allows users to create web apps through a web interface. Joomla! supports different databases, e.g., MySQL [27] and PostgreSQL [31], and a recent assessment [20] has shown that versions ranging from 3.2 to 3.4.4 suffer from a SQLi vulnerability [8].

The execution of a state-of-the-art scanner such as sqlmap on Joomla! can correctly find the vulnerability. However, sqlmap (or any other scanner for SQLi) cannot tell how that SQLi can be usefully exploited in order to carry out a concrete attack. A general description of the consequences of a SQLi attack is given in [9,29] but, whenever a SQLi entry point is found, the penetration tester has to manually investigate the kind of damages that SQLi might cause to the web app. The researchers who discovered the vulnerability of Joomla! [20] also described how it could be exploited in a real attack: it would allow an intruder to perform a session hijack and thus steal someone's session but would not allow him to create his own account or modify arbitrary data on the database. The exploration of different attack scenarios has been entirely performed manually since no automatic tool shows the outcome of the exploitation of a SQLi vulnerability on a specific web app. But who guarantees that a post-SQLi attack can actually be performed and that all possible attacks based on the SQLi have been taken into account by the penetration tester?

This is why we can't stop at the identification of a SQLi and why we can't address the post-SQLi attacks with a manual analysis. Our approach addresses this by automating the identification of attacks that exploit a SQLi.

3 SQL injections

Some general classifications based on the payloads of the SQLi (and the exploitation scenarios) have been put forth, e.g., [16,28]. Based on these, we can divide SQLi techniques into 6 different categories: (i) Boolean-Based, (ii) Time-Based, (iii) Error-Based, (iv) UNION Query, (v) Second-Order and (vi) Stacked Queries.

Given that our formalization strictly depends on the attack that the intruder wants to perform by using a particular type of SQLi, we now define the two attacks that we have considered:⁴

⁴ Other possible attacks (e.g., by exploiting a *Cross-Site Scripting* (XSS) inside the payload of some SQLi) are outside the scope of our approach for now, cf. § 7.

- *Authentication bypass attack*: the intruder bypasses an authentication check that a web app performs by querying a database.
- *Data extraction attack*: the intruder obtains data from the database that he should not be able to obtain.

Based on these attack definitions, we will now describe the main details of each category, emphasizing those aspects that are relevant for our formalization. The following table summarizes which attacks can be exploited by a SQLi technique on a specific type of SQL query. Three remarks are in order: (1) since all state-of-the-art DBMS are vulnerable to SQLi, we won't distinguish between different dialects of SQL and simply write "SQL query"; (2) for brevity, in the table we write **AB** for *authentication bypass* and **DE** for *data extraction*; (3) a scenario in which the intruder extracts information in order to bypass an authentication is considered to be a data extraction attack.

	BB		TB		EB		UQ		SO		SQ	
	AB	DE	AB	DE	AB	DE	AB	DE	AB	DE	AB	DE
SELECT	✓	✓	✓	✓	✓	✓	✓	✓				✓
UPDATE	✓	✓	✓	✓	✓	✓			✓	✓		✓
DELETE						✓						✓
INSERT						✓			✓	✓		✓

3.1 Boolean-Based SQLi (BB)

In a BB, an intruder inserts into an HTTP parameter, which is used by a web app to write a SQL query, one or more valid SQL statements that make the **WHERE** clause of the SQL query evaluate to true or false. By interacting with the web app and comparing the responses, the intruder can understand whether or not the injection was successful. In this way, an intruder can perform both authentication bypass and data extraction attacks.

In an authentication bypass attack, the intruder injects a statement that changes the truth value of a **WHERE** clause in a SQL **SELECT** query, creating a tautology. If a web app performs an authentication check querying a database, this attack will then trick the database into replying in an affirmative way even when no (or wrong) authentication details have been presented by the intruder.

In a data extraction attack, the intruder obtains data from the database. The term "extraction" is used in standard terminology but it can be misleading. With a BB, an intruder exploits the "Boolean behavior" of a web app inferring whether the original query returned some tuples or not. When the intruder understands how the web app behaves when some tuples or no tuples are returned, he can start the "extraction". In this case, the intruder asks whether a certain information is stored in the database and, based on the behavior of the web app, he knows if the information is actually inside the database.

As an example, consider a web app that presents a page *p1* containing a form. After the submission of the form, the web app creates a query for the database searching for an entry that matches the instantiation of the form fields submitted. Assume that the web app replies with *p2* if a tuple is returned, *p1* otherwise. If the intruder injects a payload such as **or username=admin** in one

of the fields of the form and is redirected to *p2*, then he will know that `admin` is a valid value of the database.

3.2 Time-Based SQLi (TB)

TB is quite similar to BB: the only difference is that TB does not need the web app to have a Boolean behavior. The intruder appends a timing function to the validity value of a Boolean clause. Thus, after the submission of the query by the web app, the database waits for a predefined amount of time for a tuple as a response to the query; the intruder can then infer whether the Boolean value of the query was true or false observing a delay in the response. As an example, consider a web app that replies with a page *p1* independent of whether a tuple is returned by the database (e.g., a search page). The intruder can inject a payload like `or if(username=admin) wait 60s` and, if `admin` is a valid entry for the column `username`, the intruder will observe a delay of around 60 seconds before receiving any answer. In real case scenarios, a BB is preferable as it is faster than a TB. Timing is not part of our formalization (see § 4), so the abstract attack traces generated by our tool will not distinguish between BB and TB.

3.3 Error-Based SQLi (EB)

When error pages are exposed to the Internet, some error messages of the database could be exposed, thus giving an intruder the possibility of exploiting an EB. In this type of injection, the intruder tricks the database into performing operations that result in an error and then he extracts information from the error messages produced by the database. EB is generally used to perform a data extraction attack by inducing the generation of an error that contains some information stored in the database. As an example, consider an intruder who wants to find out the first username in the table `usernames`. He can inject, in a login form, a payload that tricks the web app into evaluating the query `SELECT * FROM (SELECT username FROM usernames LIMIT 1) AS tbl1`. The web app generates `Error: table adminUsername unknown` because an invalid table is selected (resulting from the inner query `SELECT username FROM usernames LIMIT 1`), where `adminUsername` is the first username found in the user's table.

3.4 UNION Query-Based SQLi (UQ)

UQ is a technique in which an intruder injects a SQL UNION operator to join the original query with a malicious one. The aim is to overwrite the values of the original query and thus, in order to extract information, UQ requires the web app to print the result of the query within the returned HTML page. This behavior allows the intruder to actually extract information from the database by reading it within the web app itself. As an example, consider the query `SELECT nickname FROM users WHERE id=$id`. An intruder can inject `$1 UNION ALL SELECT creditCardNumber FROM CreditCardTable` as `id` and thus obtain all the credit card numbers.

3.5 Second-Order SQLi (SO)

SO is an injection that has no direct effect when submitted but that is exploited in a second stage of the attack. In some cases, a web app may correctly handle and store a SQL statement whose value depends on the user input. Afterwards, another part of the web app that doesn't implement a control against SQLi might use the previously stored SQL statement to execute a different query and thus expose the web app to a SQLi. Automated scanners generally fail to detect this type of SQLi (e.g., [34,35]) and may need to be manually instructed to check for evidence that an injection has been attempted. As an example, consider a web app in which there is a user with username `admin` (stored into the database). If this web app has a registration page that allows special characters in the username, an intruder can register a new user with username `admin'#` (where we assume that `#` is the comment delimiter character) and log in as `admin'#`. If the intruder changes the password of the user `admin'#`, a query like `UPDATE users SET password='123' WHERE username='admin'#` is executed. The *Database Management System (DBMS)* will interpret the `WHERE` clause as `username='admin'` (because everything after the `#` is considered a comment), so that the intruder succeeds in changing the password of the admin user.

3.6 Stacked Queries SQLi (SQ)

With a SQ,

an intruder can execute an arbitrary query different from the original one. The semicolon character `;` enables the intruder to concatenate a different SQL query to the original one. By doing so, the intruder can perform data extraction attacks as well as execute whatever operation is allowed by the database. With a SQ, an intruder can perform any of the SQLis described above. Thus, whenever we refer to all the SQLis in our categorization, we exclude SQ as it is already covered by the other ones.

Prevention techniques. Avoiding SQLi attacks is theoretically quite straightforward. In fact, developers can use sanitization functions or prepared statements. Roughly speaking, the general idea is to not evaluate the injected string as a SQL command.

A *sanitization function* takes the input provided by the user and removes (i.e., escapes) all the special characters that could be used to perform a SQLi. Sanitization functions are not the best option when dealing with SQLi because they might not be properly implemented or do not consider some cases.

Prepared statements are the best option for preventing SQLis. They are mainly used to execute the same query repeatedly maintaining efficiency. However, due to their inner execution principle (if properly implemented) they are immune to SQLi attacks. The execution of a prepared statement consists mainly in two steps: preparation and execution. In the preparation step, the query is evaluated and compiled, waiting for the parameters for the instantiation. During the execution step, the parameters are submitted to the prepared statement and handled as data and thus they cannot be interpreted as SQL commands.

4 A Formalization of SQLi

We will now describe how we formally represent a web app that interacts with a database using insecure SQL queries and/or a sanitized (i.e., secure) query. In § 4.1, we propose an extension of the DY model that can deal with SQLi.⁵ We formalize the database in § 4.2, the web app in § 4.3, and the goals in § 4.4. For brevity and readability, we omit many details and only give pseudo-code that should be quite intuitive. See the appendix for full details and the ASLan++ code of our formalizations and case studies, along with a brief introduction to ASLan++.

4.1 The DY web intruder

We extend the standard DY intruder model [11] for security protocol analysis. Suppose that we want to search for an authentication bypass attack via BB (§ 3), in which the intruder injects a statement that changes the truth value of a `WHERE` clause in a SQL `SELECT` query, creating a tautology. To formalize this, we need to extend the DY intruder by giving him the ability to send a concatenation of Boolean formulas made of conjunctions and disjunctions. This characteristic highlights an important difference between the classical DY intruder and the enhanced version we are proposing: *our web intruder works with abstract payloads rather than messages*. Due to technical details (e.g., implementation constraints and non-termination problems), implementing such a modification is impractical. We have thus allowed the intruder to concatenate the exact payload, *or true*, and defined a Horn clause to model that whenever a formula has *or true* injected by the intruder, it evaluates to true.

We can rephrase the same reasoning in the case of BB for data extraction attacks, in which the intruder tricks the web app into asking to the database if a particular information is present; for example, instead of `or true`, the intruder adds `or username=admin`. The DBMS will reply in an affirmative way only if there is a tuple in the database with `admin` as username. To allow the intruder to perform all the SQLis described in § 3, we thus extend the DY intruder with one constant `sqli` that represents any SQLi payload (e.g., `or true`).

4.2 The database

We give a general formalization of a database that can be used in any specification to exploit SQLi when searching for security flaws in a web app. Our formalization aims to be both *compact*, to avoid state-space explosion problems, and *general enough* not to be tailored to a given technology (e.g., MySQL

⁵ This formal representation is intended to work with tools that perform symbolic analysis. We don't formalize the honest client behavior and we assume the DY intruder to be the only agent able to communicate with the web app. The DY intruder will eventually perform honest interactions if needed to achieve a particular configuration of the system.

or PostgreSQL). Hence, we don't represent the database content, the database structure, the SQL syntax nor access policies specified by the DBMS. Rather, we formalize messages sent and received and queries, and a database can be seen as a network node that interacts only with the web app through a secure channel.⁶

Definition 1. Messages consist of variables V , constants c (*sqli*, etc.), concatenation $M.M$, function application $f(M)$ of uninterpreted function symbols f to messages M (e.g., $\mathbf{tuple}(M)$), and encryption $\{M\}_M$ of messages with public, private or symmetric keys that are themselves messages.⁷ We define that M_1 is a submessage of M_2 as is standard (e.g., M_1 is a submessage of $M_1.M_3$, of $f(M_1)$ and of $\{M_1\}_{M_4}$) and, abusing notation, write $M_1 \in M_2$.

Definition 2. A query is valid (respectively, not valid) when, evaluated by a database, it returns one or more (respectively, zero) tuples.

We formalize the validity of SQL queries by means of the Horn clause: $\mathbf{inDB}(M.\mathbf{sqli}) \implies \mathbf{true}$, where, in order to represent a SQLi attack, the predicate $\mathbf{inDB}()$ holds for a message (which represents a SQL query) whenever it is of the form $M.\mathbf{sqli}$. This states that the intruder has injected a payload *sqli* into the query parameters (expressed as a variable) M .

Incoming messages. We consider, as incoming messages, only SQL queries via raw SQL and via sanitized queries. The parameters of queries are represented by a generic variable *SQLquery*. In case of a raw SQL query, they are wrapped by an uninterpreted function *query()*; if a sanitized query has been implemented then we use another uninterpreted function *sanitizedQuery()*. These two uninterpreted functions allow the modeler to “switch on/off” the possibility of a SQLi in some point of the app.

Database responses. The tuple generated by the database as a response to a raw SQL query is represented by an uninterpreted function *tuple()* over a message representing a SQL query. Given that we do not model the content of the database, this function represents any (and all) database data.

Whenever the database receives a SQL query *query(SQLquery)* from the web app, the uninterpreted function *tuple(SQLquery)* is sent back to the web app to express that a tuple, as a response to the query, has been found. This response is returned only if $\mathbf{inDB}()$ holds; in all other cases, a constant *no_tuple* is returned to represent that no tuples are returned in the responses of the database.

If the database receives a sanitized query, no injection is possible. Hence, the database does not return any useful information to the web app; instead, a constant *no_tuple* is returned. Since the intruder cannot perform a SQLi in presence of a sanitized query, we also assume that a sanitized query can be

⁶ Nothing prevents us from relaxing this assumption but this would give the DY intruder the possibility of performing attacks (e.g., man-in-the-middle attacks) that are rare in web app scenarios.

⁷ In this paper we do not need to distinguish between the different kinds of encrypted messages, but we could of course do by following standard practice.

executed only with legitimate parameters, i.e., as a function of `tuple()` (this is because we are interested in modeling only SQLi scenarios).

The pseudo-code representing the database behavior is given in Listing 1.1, where, here and in the following, we write DB for the database. DB is a network node and we assume it to be always actively listening for incoming messages. It is defined by two main, mutually exclusive, branches of an if-elseif statement: one guard is in line 1 in which DB is waiting (expressed in Alice-and-Bob notation) for a sanitized query and the other in line 3 in which it is waiting for a raw SQL query. If a sanitized query is received, then there is no SQLi. Given that we only consider dishonest interactions, the data sent back to the intruder will not increase his knowledge. In other words, no SQLis are permitted and any permitted query will just give to the intruder the possibility of continuing his execution with the web app but won't add any extra information to his knowledge.

Listing 1.1. Pseudo-code of a DBMS.

```

1  if(WebApp -> DB: sanitizedQuery(SQLQuery)){
2    if(SQLQuery == tuple(*)) DB -> WebApp: no_tuple;
3  }elseif(WebApp -> DB: query(SQLQuery)){
4    if(inDB(SQLQuery)) DB -> WebApp: tuple(SQLQuery);
5    if(!(inDB(SQLQuery))) DB -> WebApp: no_tuple;}

```

One may argue that a valid query should indeed add extra information to the intruder knowledge. However, we do not model the content of the database and any information received by the intruder as a response to a sanitized query is included in the action that the web app performs after this database response. Thus, in our formalization, the query in `SQLQuery` is not valid and then the `no_tuple` constant is sent back in line 2. We also add a constraint in line 2 that any query received (`SQLQuery`) must be of the form `tuple(*)`, i.e., as a function of the content of the database where `*` acts as a wildcard character that matches any possible parameter. This is because, in the case of a sanitized query, the intruder cannot perform a SQLi and we exclude the case in which the DY intruder sends a random query just to continue the execution with the web app. Instead, he has to either know a tuple of the database or data as functions of a tuple of the database. In the case the intruder knows `tuple(Query)`, he will just receive `no_tuple`, i.e., correctly no data has been leaked to the intruder.

The second branch of the initial if-elseif statement (line 3) handles raw queries. If a raw query is submitted, then there are two cases: the raw query is not valid (line 5, where `!` formalizes the negation) and then, as in the previous case, `no_tuple` is sent back (line 5); the raw query is valid (line 4) and a tuple is sent back (line 4). Given that all these queries are sent from the intruder, we can assume they have a malicious intent. One may argue that, in a real case scenario, the database is not actually returning a tuple but, given that an intruder could repeatedly send a SQLi exploiting that injection point, it is fair to assume that the database is sending all the tuples it contains, i.e., `tuple(SQLQuery)`.

4.3 The web app

As for the database, the web app is a node of the network that can send and receive messages. The web app communicates with a client or with the database

(it can, potentially, also communicate with other apps but we do not consider that explicitly here). We assume only one database is present because adding other databases would not add any further useful information for finding attacks based on SQLi. The proof is straightforward. Since we do not consider database contents and structures, if we wanted to have two database models, then we would have two exact copies of the formalization given in § 4.2. Since we have assumed that there exists a long-lasting secure relation between the database and the web app, no man-in-the-middle attacks are considered. Therefore, any attack found that involves the communication with one of the two databases could be found by considering the other database only.

A specification of a web app can be seen as a behavioral description of the web app itself (along with its interaction with the database). A modeler can define this specification from the design phase documentation of the engineering process of the web app. A model can also be created in a black box way by just looking at the HTTP messages exchanged from a client and the web app and guessing the communication with the database.

We now consider the main aspects that allow for the modeling of a web app.

Sending and receiving messages. A web app can communicate with a client and the database. We abstract away as many details as possible of the web pages and thus any incoming message will only contain: (i) parameters of forms expressed as variables, e.g., `Client -> WebApp: Username.Password`, and (ii) the web page itself expressed as a constant, e.g., `WebApp -> Client: dashboard` where `dashboard` represents a web page. Note that, in any response of the web app, if the content of the response is linked to a response of the database, i.e., `tuple(Query)` (where the query is either `SELECT`, `UPDATE` or `DELETE`), then `tuple(Query)` must be included in the response. Otherwise, we would end up representing a scenario in which no content of the tuple received by the database is included in, or linked to, the web page and thus no SQLi would be present.

Queries. The web app creates either a sanitized query or a raw query. Then, the web app wraps the variables representing the query parameters with either `sanitizedQuery()` or `query()`, both uninterpreted functions, and afterwards sends the SQL query to the database. Note that we only need to represent the parameters of a SQL query since we do not distinguish between different queries in the database formalization. If the query does not depend on parameters sent from a client, the intruder cannot exploit it to perform a SQLi. The SQL query used to query the database is represented as a constant, resulting in the database always replying with `no_tuple` (as `inDB()`, in this case, is never valid).

If statements. We use them mainly to decide, based on which kind of message has been received, what the web app has to reply. For example, if the database replies with a tuple `tuple(Query)`, then the web app might return a specific page along with `tuple(Query)` or might return a different page.

Assignments. A constant or a message can be assigned to a variable: `Variable := constant|message`. Assignments are, e.g., useful to save incoming messages.

4.4 Goals

Finally, we define the security properties we want the model to satisfy. As we discussed in § 3, we consider two main attacks: authentication bypass and data extraction. We give the formalization in Listing 1.2, where `iknowledge` is a predicate that represents the knowledge of the intruder. By using the LTL “globally” operator `[]`, we can specify an authentication check by stating that the intruder should not have access to a specific page (`dashboard` in Listing 1.2), whereas data extraction is represented by specifying that the intruder should not increase his knowledge with data from the database (i.e., as function of `tuple()`).

Listing 1.2. Authentication bypass and data extraction goals of the BB example.

```
[](! (iknowledge(dashboard))); %authentication bypass
>[](! (iknowledge(tuple(*)))); %data extraction
```

5 SQLfast, Case Studies and Results

To show that our formalization can be used effectively to detect security flaws linked to SQLi attacks, we have developed *SQLfast*, a prototype *SQL Formal Analysis Tool* [33]. In [33] we also provide a friendly web-based user interface that helps the modeler in creating the web-app model. *SQLfast* takes in input a specification written in ASLan++ [39], the modeling language of the AVANTSSAR Platform for security protocol analysis [3], and, by using the ASLan++ translator, generates a transition system in the low-level language ASLan [4]. It then calls the model checker CL-AtSe [37] and generates an *Abstract Attack Trace (AAT)* as a MSC if an attack was found. *SQLfast* automatically detects which type of SQLi has been exploited and, in an interactive way, generates the curl or sqlmap [34] commands to concretize the attack reported.

As a concrete proof-of-concept, we have applied *SQLfast* to (i) WebGoat [30], (ii) Damn Vulnerable Web Application (DVWA) [13], (iii) Joomla! 3.4.4, and (iv) *Yet Another Vulnerable Web Application (YAVWA)*, an ad-hoc testing environment that we have developed and that also includes a SO SQLi example. The case studies provided by WebGoat and DVWA might sound limited but capture all possible scenarios with respect to SQLi attack combinations considered in this paper — recall that our formalization for SQLi attacks does not find SQLi payloads, but focuses on vulnerabilities based on SQLi. We tested *SQLfast* in order to show all the combinations that could be represented by considering SQLi for (1) authentication bypass, (2) data extraction and (3) data extraction with reuse of the extracted information. Our case studies are quite heterogenous, so it should not be difficult to map other case studies to one of these scenarios we have considered.

We have implemented the case studies in ASLan++ to be able to apply the model checkers of the AVANTSSAR Platform (in particular, CL-AtSe), but other model checkers implementing the Dolev-Yao intruder model could be used as well, provided that their input language is expressive enough. For the sake of brevity, we discuss here only the case studies Joomla!, YAVWA and SO, which

show how our formalization can find attacks linked to the logic of a web app that is vulnerable to SQLi attacks. The type of attacks that SQLfast can detect and concretize cannot be detected by state-of-the-art tools for SQLi such as sqlmap.

5.1 Case Study: Authentication Bypass via Data Extraction

We now discuss two scenarios in which our approach detects attacks that state-of-the-art tools, such as sqlmap, are *not* able to detect and exploit. In the first scenario, the intruder exploits a recent SQLi vulnerability found (by manual inspection only) in Joomla! [8]. The second scenario (YAVWA) is a variant of the first and shows a concatenation of different attacks.

Joomla!. A recent assessment has shown that the Content History module of Joomla! suffers from a SQLi vulnerability that allows a remote (non-authenticated) user to execute arbitrary SQL commands [8]. The pseudo-code in Listing 1.3 represents the following behavior: a remote user visits the Content History component (line 1). The web app queries the database with the user supplied data (2). If some tuples are generated (3), the web app sends to the client the history page `viewHistory` along with the `tuple()` function (4). The web app then has two possible ways of authenticating the user (5–9): by using credentials or cookies. If username and password are provided (5), the web app applies a non-invertible hash function `hash()` to the password, and queries the database to verify the credentials (6).⁸ If the credentials are correct, the administration panel is sent to the user (7). In case of a cookie session, the user provides a cookie that the web app checks querying the database (8). If the cookie is valid, the administration panel is sent back to the user (9).

Listing 1.3. Pseudo-code representing the Joomla! scenario.

```

1 User -> WebApp: com_contenthistory.history.Listselect;
2 WebApp->DB: query(com_contenthistory.history.Listselect);
3 if(DB -> WebApp: tuple(SQLquery)){
4   WebApp -> User: viewHistory.tuple(SQLquery); }
5 if(User -> WebApp: Username.Password){
6   WebApp -> DB: sanitizedQuery(Username.hash>Password));
7   if(DB -> WebApp: no_tuple){ WebApp -> User: adminPanel; }}
8 if(User -> WebApp: Cookie){ WebApp -> DB: sanitizedQuery(Cookie);
9   if(DB -> WebApp: no_tuple){ WebApp -> User: adminPanel; }}
```

As goal, we check if there exists an execution in which the intruder can access the administration panel represented by the constant `adminPanel`.

Listing 1.4. Authentication bypass for the Joomla! scenario.

```
[](!(iknowledge(adminPanel)));
```

SQLfast generates the AAT in Listing 1.5, which is an authentication bypass attack where the intruder hijacks a user session by using a cookie instead of login credentials. In fact, the web app applies a hash function to the password before verifying the credentials submitted by the user. The hash function would

⁸ The web app applies a hash function to the password before checking whether credentials are correct because Joomla! stores the passwords hashed into the database.

not allow an intruder to blindly submit a password extracted from the database, the only possibility is using a valid cookie value.⁹ The intruder performs a data extraction and retrieves the information to access the administration panel (1–4), and uses it to hijack a user session by submitting a valid cookie value (5–8).

Listing 1.5. Abstract attack trace that extracts data with a SQLi in order to bypass the authentication of the Joomla! scenario.

```

1 i -> WebApp : com_contenthistory.history.sqli
2 WebApp -> DB : query(com_contenthistory.history.sqli)
3 DB -> WebApp : tuple(com_contenthistory.history.sqli)
4 WebApp -> i : viewHistory.tuple(com_contenthistory.history.sqli)
5 i -> WebApp : cookie.tuple(com_contenthistory.history.sqli)
6 WebApp -> DB : sanitizedQuery(tuple(com_contenthistory.history.sqli))
7 DB -> WebApp : no_tuple
8 WebApp -> i : adminPanel

```

YAVWA. We have designed a variant of Joomla! to show that a SQLi can be exploited to compromise a part of a web app that does not directly depend on databases. YAVWA provides an HTTP form login and a login by HTTP basic authentication [19] configured with the `.htaccess` [2] file. The credentials used for the HTTP basic authentication, which are stored in the `.htpasswd` file, are the same as the ones employed by the users to login into the web app (i.e., the same as the ones stored in the database).¹⁰ The intruder’s goal is to access the area protected by the HTTP basic authentication login. Obviously, he cannot perform a SQLi to bypass HTTP basic authentication since the login procedure doesn’t use SQL. Bypassing the login page, without knowing the correct credentials, doesn’t allow the intruder to gain access to the secure folder.

We have defined this scenario in the pseudo-code in Listing 1.6. The client sends his personal credentials (`Username.Password`) to the web app (1). The web app creates a query that it sends to the database (2) for verifying the submitted credentials. If tuples are generated from the database (3), a dashboard page is returned to the client along with the function `tuple()` (4), otherwise, the web app redirects the user to the login page (5). At this point, the web app waits to receive correct credentials that will allow the client to access the secure folder `secureFolder` (6). Given that the credentials are the same as the ones stored in the database, and the database content is represented with the function `tuple()`, we can also represent credentials here with the function `tuple()`.¹¹

Listing 1.6. Pseudo-code representing the YAVWA scenario.

```

1 User -> WebApp: Username.Password;
2 WebApp -> DB: query(Username.Password);

```

⁹ We do not consider the possibility of brute forcing the hashed password, in accordance with the perfect cryptography assumption of the DY model.

¹⁰ One might see this as a bad practice and thus consider this whole example to be unrealistic. However, this scenario can easily be categorized as a security misconfiguration (5th most critical security issue according to [28]). Moreover, recent events [15] have shown, once again, how humans tend to reuse passwords across multiple logins.

¹¹ We recall from § 4.2 that `tuple()` represents an abstraction of any data that can be extracted from the database. This means that whenever a web app requires any data in the domain of the database, we can write them as a function of `tuple()`.

```

3  if(DB -> WebApp: tuple(SQLQuery)){
4    WebApp -> User: dashboard.tuple(SQLQuery);
5  }elseif(DB -> WebApp: no_tuple){ WebApp -> User: login; }
6  if(User -> WebApp: tuple(*)){ WebApp -> User: secureFolder; }

```

As goal, we check if the intruder can reach **secureFolder**. SQLfast generates the AAT given in Listing 1.7, in which the intruder successfully retrieves information from the database and uses such information to access a protected folder. The intruder performs a data extraction attack using SQLi (1–4), which allows him to retrieve information stored in the database, and then (5–6) submits the extracted data and accesses the restricted folder **secureFolder**.

Listing 1.7. Abstract attack trace of the YAVWA case study.

```

1  User -> WebApp: Username(4).sqli
2  WebApp -> DB : query(Username(4).sqli)
3  DB -> WebApp : tuple(Username(4).sqli)
4  WebApp -> i : dashboard.tuple(Username(4).sqli)
5  i -> WebApp : tuple(Username(4).sqli)
6  WebApp -> i : secureFolder

```

5.2 Case Study: Second-Order SQLi (SO)

We now show that our formalization is flexible enough to represent SOs, which are notoriously very difficult to detect and exploit.

This scenario is part of YAVWA and implements a web app that allows users to register a new account. In the registration process, the web app executes an (INSERT) SQL query that stores the user’s credentials into a database. The intruder can create an account submitting malicious credentials that don’t result in a SQLi but will trigger an injection later on in the web app. After the registration phase, the user submits a request for accessing an internal page. The web app performs another SQL query using the same parameters previously used in the registration process (i.e., the registration credentials). At this point, a page is showed together with the injection and the intruder can exploit a SO.

We have formalized this scenario in Listing 1.8: a client sends a registration request along with his personal credentials (**Username** and **Password**) to the web app (1). The web app sends a query containing the client’s credentials to the database (2). The web app checks if it receives a response from the database containing the data resulting from the execution of the query **tuple(SQLQuery)** submitted by the web app (3). The web app sends back to the client the page **registered** (4). Here, the web app does not forward **tuple()** because the registration query is an **INSERT** (see § 4.2). The client asks for a page (5), which makes the web app use previously submitted values of **Username** and **Password** to execute a new SQL query (6). Here is where the SO takes place; the variables embedded in the query in (6) will trigger a SO. The database executes the query and sends back the results to the web app (7). Finally (8), the web app sends (by using a **SELECT** query) to the client the requested page and the **tuple()**.¹²

¹² Recall that we don’t represent SQL syntax in our models, so we don’t explicitly represent the type of the SQL according to the modeling guidelines in § 4.3.

Listing 1.8. Pseudo-code representing a web app vulnerable to a SO attack.

```
1 User -> WebApp: registrationRequest.Username.Password;
2 WebApp -> DB: query(Username.Password);
3 if(DB -> WebApp: tuple(SQLQuery)){
4   WebApp -> User: registered;
5   User -> WebApp: requestPage;
6   WebApp -> DB: query(Username.Password);
7   DB -> WebApp: tuple(SQLQuery);
8   WebApp -> User: page.tuple(SQLQuery); }
```

As goal, we ask if the intruder can interact with the web app until he obtains data from the database, i.e., with a data extraction attack, as in Listing 1.2. SQLfast generates the AAT in Listing 1.9, in which the intruder performs the registration process (1–4) by registering malicious credentials `Username(4)` and `sqli`. At the end of the registration process (5), the intruder asks for `requestPage` that makes the web app send to the database a SQL query with the same parameters the intruder used in the registration (6–7). In (8), the intruder receives the requested page and the result of the execution of the injected SQL query performing a SO.

Listing 1.9. Abstract attack trace for the SO case study.

```
1 User -> WebApp: registrationRequest.Username(4).sqli
2 WebApp -> DB : query(Username(4).sqli)
3 DB -> WebApp : tuple(Username(4).sqli)
4 WebApp -> i : registered
5 i -> WebApp : requestPage
6 WebApp -> DB : query(Username(4).sqli)
7 DB -> WebApp : tuple(Username(4).sqli)
8 WebApp -> i : page.tuple(Username(4).sqli)
```

5.3 Concretization phase

We executed SQLfast on all our case studies using a standard laptop (Intel i7 with 8G RAM). The execution time of the model-checking phase of SQLfast ranges from 35 to 45 ms. The overall process (from translation to concretization) takes a few seconds. In all the cases, we generated AATs violating the security property we defined over the model (authentication bypass or data extraction attack). Once the AAT has been generated, SQLfast interactively asks the user to provide information such as the URL of the web app. Finally, if we are concretizing a SQLi that exploits an authentication bypass attack a curl command is showed, whereas sqlmap is used for data extraction SQLi. By executing the traces generated by SQLfast, we exploited all the AATs over the real web app.

6 Related work

Many works have proposed new SQLi techniques and payloads (e.g., [36,22,10]) or formal approaches to detect SQLi (e.g., [26,24,17,25]). However, to the best of our knowledge, ours is the first attempt to search for vulnerabilities based on SQLi rather than to detect SQLi. There are, however, a number of works that are closely related to ours and that are thus worth discussing.

SPaCiTE is a model-based security testing tool for web apps that relies on mutation testing [5]. SPaCiTE starts from a secure ASLan++ specification of

a web app and automatically introduces flaws by mutating the specification. The strength of this approach is the concretization phase. Starting from an AAT, generated from the mutated specification using a model-checking phase, SPaCiTE concretizes and tests the attack trace on the real web app. The major differences with respect to our approach reside in how we model web apps and in particular those aspects that strictly characterize SQLi aspects. The main goal of the approach in [5] is to find SQLi entry points and concretize them, our main goal is to consider SQLi aspects that can be exploited to attack a web app.

Another formal approach that uses ASLan++ and the DY intruder model for the security analysis of web apps is [32]. In this work, the authors model a web app searching for CSRF and they do not consider databases or extensions to the DY model. However, the idea and the representation of web apps is close to ours and we envision some potentially useful interaction between the two approaches.

In [6], the authors describe the “Chained Attack” approach, which considers multiple attacks to compromise a web app. The idea is close to ours, but: (i) they consider a new kind of web intruder, whereas we stick with the DY intruder; (ii) we analyzed the most common SQLi techniques and proposed a formalization of a vulnerable database, they only consider the behavior of the web app.

In [1], the authors present a model-based method for the security verification of web apps. They propose a methodology for modeling web apps and model 5 case studies in Alloy [21]. Even if the idea is similar to our approach, they have defined three different intruder models that should find web attacks, whereas we have used (and extended) the standard DY one. Their AATs are difficult to interpret because no MSCs are given but state configurations. They have also considered a number of HTTP details that we have instead abstracted away in favor of an easier modeling phase. In contrast, we display AAT as MSCs and we proposed a concretization phase to obtain the concrete payloads of SQLi.

7 Conclusions and future work

We have presented a formal approach for the representation of SQLi and attacks that exploit SQLi in order to violate security properties of web apps. We have formally defined web apps that interact with a database (that properly replies to queries containing SQLi) and an extended DY intruder able to deal with authentication bypass and data extraction attacks related to SQLi. We have shown the efficiency of our prototype tool SQLfast on four real-world case studies. SQLfast handles SO and detects multi-stage attacks and logical flaws that, to the best of our knowledge, no other tool can handle together, and hardly ever even individually, including the discovery of an attack on Joomla!.

As future work, we plan to extend the database formalization in order to consider SQLi that would modify the database state leading to more complex SQLi exploitations. We also plan to analyze other web app vulnerabilities such as stored/reflected XSS and broken session management, and investigate synergies between our approach and the one of [32] on CSRF. We will extend our approach to detect (i) complex concatenations of vulnerabilities (similar to, and more

complex than, [18]) that lead to concatenations of attacks, and (ii) articulated paths to vulnerabilities that would hardly ever be discovered by manual analysis.

References

1. D. Akhawe, A. Barth, P. Lam, J. Mitchell, and D. Song. Towards a Formal Foundation of Web Security. In *CSF*, pages 290–304. IEEE, 2010.
2. Apache software foundation. Apache HTTP Server Tutorial: .htaccess files. <https://httpd.apache.org/docs/current/howto/htaccess.html>.
3. A. Armando et al. The AVANTSSAR Platform for the Automated Validation of Trust and Security of Service-Oriented Architectures. In *TACAS 2012*, LNCS 7214, pages 267–282. Springer, 2012.
4. AVANTSSAR. Deliverable 2.3 (update): ASLan++ specification and tutorial, 2011. <http://www.avantssar.eu>.
5. M. Büchler, J. Oudinet, and A. Pretschner. Semi-automatic security testing of web applications from a secure model. In *SERE*, pages 253–262, 2012.
6. A. Calvi and L. Viganò. An Automated Approach for Testing the Security of Web Applications Against Chained Attacks. In *ACM/SIGAPP SAC*. ACM Press, 2016.
7. S. Christey. The 2009 CWE/SANS Top 25 Most Dangerous Programming Errors. <http://cwe.mitre.org/top25>.
8. CVE-2015-7857. <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2015-7857>.
9. CWE. CWE-89: Improper Neutralization of Special Elements used in an SQL Command (‘SQL Injection’). <https://cwe.mitre.org/data/definitions/89.html>.
10. B. Damele and A. Guimarães. Advanced SQL injection to operating system full control. In *BlackHat EU*, 2009.
11. D. Dolev and A. C. Yao. On the security of public key protocols. *IEEE Trans. Inf. Theory*, 29(2):198–208, 1983.
12. A. Doupé, M. Cova, and G. Vigna. Why Johnny Can’t Pentest: An Analysis of Black-Box Web Vulnerability Scanners. In *DIMVA*, LNCS 6201, pages 111–131. Springer, 2010.
13. Damn Vulnerable Web Application (DVWA). <http://www.dvwa.co.uk>.
14. J. Forristal. ODBC and MS SQL server 6.5. *Phrack*, 8(54), 1998.
15. The Hacker News. <http://thehackernews.com/2016/06/oculus-brendan-iribe.html>.
16. W. G. Halfond, J. Viegas, and A. Orso. A Classification of SQL-Injection Attacks and Countermeasures. In *SIGSOFT ’06/FSE-14*, 2006.
17. W. G. J. Halfond and A. Orso. AMNESIA: Analysis and Monitoring for NEutralizing SQL-injection Attacks. In *ASE*, pages 174–183. IEEE, 2005.
18. E. Homakov. How I hacked Github again, 2014. <http://homakov.blogspot.it/2014/02/how-i-hacked-github-again.html>.
19. Internet Engineering Task Force (IETF). HTTP Authentication: Basic and Digest Access Authentication, 1999. <https://www.ietf.org/rfc/rfc2617.txt>.
20. iSpiderLabs. Joomla SQL Injection Vulnerability Exploit Results in Full Administrative Access, 2015. <https://www.trustwave.com/Resources/SpiderLabs-Blog/Joomla-SQL-Injection-Vulnerability-Exploit-Results-in-Full-Administrative-Access/?page=1&year=0&month=0>.
21. D. Jackson. *Software Abstractions: Logic, Language, and Analysis*. MIT Pr., 2012.
22. O. M. Jayathissa. SQL Injection in Insert, Update and Delete Statements.

23. Joomla! <https://www.joomla.org>.
24. A. Kiezun, P. J. Guo, K. Jayaraman, and M. D. Ernst. Automatic creation of SQL injection and cross-site scripting attacks. In *ICSE*, pages 199–209. IEEE, 2009.
25. V. B. Livshits and M. S. Lam. Finding Security Vulnerabilities in Java Applications with Static Analysis. In *USENIX*, pages 18–18, 2005.
26. M. Martin and M. S. Lam. Automatic Generation of XSS and SQL Injection Attacks with Goal-Directed Model Checking. In *USENIX*, pages 31–43, 2008.
27. MySQL. <https://www.mysql.com>.
28. OWASP. Owasp top 10 for 2013. https://www.owasp.org/index.php/Category:OWASP_Top_Ten_Project.
29. OWASP. SQL Injection. https://www.owasp.org/index.php/SQL_Injection.
30. OWASP. WebGoat Project. https://www.owasp.org/index.php/Category:OWASP_WebGoat_Project.
31. PostgreSQL. <http://www.postgresql.org>.
32. M. Rocchetto, M. Ochoa, and M. Torabi Dashti. Model-Based Detection of CSRF. In *IFIP SEC*, pages 30–43. Springer, 2014.
33. SQLfast: SQL Formal Analysis Tool, 2015. <http://sqlfast.altervista.com>.
34. sqlmap: Automatic SQL injection and database takeover tool, 2013. <http://sqlmap.org>.
35. sqlninja: a SQL Server injection & takeover tool, 2013. <http://sqlninja.sourceforge.net>.
36. M. Stampar. Data Retrieval over DNS in SQL Injection Attacks. <http://arxiv.org/abs/1303.3047>, 2013.
37. M. Turuani. The CL-AtSe Protocol Analyser. In *RTA*, LNCS 4098, pages 277–286, 2006.
38. L. Viganò. The SPaCIoS Project: Secure Provision and Consumption in the Internet of Services. In *ICST*, pages 497–498, 2013.
39. D. von Oheimb and S. Mödersheim. ASLan++ — a formal security specification language for distributed systems. In *FMCO*, LNCS 6957, pages 1–22. Springer, 2010.

A Correctness of the database formalization

In this section, we prove that the formalization of the database correctly handles all the SQLis categorized in § 3. We remark that tool, specification and guidelines show that it is concretely possible to exploit these SQLis.

Theorem 1. *Let i represent our extension of the DY intruder (cf. § 4.1) and db represent the database. For any message Q such that $sqli$ is a submessage of Q (i.e., $sqli \in Q$), we have:*

$$\begin{aligned}
 &\text{if } i \rightarrow db : \text{sanitizedQuery}(Q) \text{ or} \\
 &\quad i \rightarrow db : \text{query}(Q) \\
 &\text{then } db \rightarrow i : \text{Response}
 \end{aligned}$$

where $i \rightarrow db : f(M)$ means that i sends to db a message M through the web app, which creates a query from the message M either using a raw query if $f = \text{query}$, and a sanitized query otherwise, and Response is (according to the formalization proposed in § 4.2) a variable such that its instantiation is either no_tuple if sanitizedQuery has been used, or $\text{tuple}(Q)$ otherwise, with $\text{tuple}(Q)$ representing the exploitation of a SQLi $s \in Q$.

Proof. We have extended the DY intruder with `sqli` to represent the payload that the intruder can use to send a SQLi to the web app and then to the database.

Given that `sqli` $\in Q$, we only have two subcases: the query Q is either (i) a sanitized query or (ii) a raw query.

- (i) Line 1 of the database (whenever in the proof we refer to a line of the database we are referring to Listing 1.1) accepts the incoming message and returns the constant `no_tuple` if Q contains *only* a tuple of the database (line 2). In this case, `SQLquery!=tuple(*)` because contains `sqli` and nothing is returned (i.e., a SQLi is detected). Correctly, no SQLi can be performed in this case.
- (ii) Line 3 of the database accepts the incoming message and, given that the message contains a SQLi, the predicate `inDB(Q.sqli)` is true (due to the Horn clause `inDB(M.sqli) \implies true`, defined in § 4.2), and the guard in line 4 is also true. The database then replies with `tuple(Q.sqli)` and a SQLi has been exploited.

□

B ASLan++ formalization

Although in the body of the paper we have shown only intuitive pseudo-code, we have actually implemented our formalizations using ASLan++. In § B.1, we give a brief overview of ASLan++ [39] and ASLan [4], which are used by the AVANTSSAR Platform [3] and the SPaCIoS Tool [38] for the formal and automated analysis at design-time and validation at run-time of security-sensitive protocols and web apps and their security goals. We focus on the aspects that we use for modeling our case studies of § 5. In § B.2, we describe an ASLan++ specification skeleton that encapsulates all the aspects considered in § 4 and then expand this to full specifications of the case studies in § C.

B.1 ASLan++ and ASLan

ASLan++ is a formal and typed security protocol specification language, whose semantics is defined in terms of ASLan, which we describe below. Similar to object-oriented languages, an ASLan++ specification consists in a hierarchy of *entity declarations*, which are similar to Java classes.

The top-level entity is usually called *Environment* (similar to the “main” procedure of a program) and it typically contains the definition of a *Session* entity, which in turn contains a number of subentities (and their instantiations, i.e., `new subentity(<parameters>);`) that define the main principals involved in the system (e.g., web app, database). Each subentity defines the internal behavior of the component it models and the interaction with other entities. We won’t go into the full details of the ASLan++ language [4], but in the following we discuss in itemized form the main aspects that we use to give the ASLan++ code of our formalizations.

- ASLan++ supports variables (capital letter symbols), constants (lower case symbols), functions and predicates that can be distinguished by the return type, message and fact (described afterwards) respectively.
 - An entity is composed by two main sections:
 - **symbols**, in which types of local symbols are declared and instantiations are given of all the variables and constants used in the entity, and
 - **body**, where the behavior of the entity is defined. The instantiation of an entity must be done in the body of the parent-entity using the **new** keyword as follows: **new subentity(<params>)**.
- Inside the body of an entity we use three different types of statements: assignments, message send and message receive. The *assignment* **Var:=fresh()** assigns to the variable **Var** a new constant of the proper type, while **Var:=Var'** is the usual assignment.
- Sending and receiving of messages are expressed in Alice-and-Bob notation: **A -> B: M;**, where **A** and **B** are entities and **M** a message. More formally, a *message send* statement, **Snd -> Rcv: M**, is composed by two variables **Snd** and **Rcv** representing sender and receiver, respectively, and a message **M** exchanged between the two parties. In *message receive*, **Snd** and **Rcv** are swapped and usually, in order to assign a value to the variable **M**, a **?** precedes the message **M**, i.e., **Snd -> Rcv: ?M**. However, in ASLan++, the **Actor** keyword refers to the entity itself (similar to “this” or “self” in object-oriented languages) and thus we actually write the send and receive statements as **Actor -> Rcv: M** and **Snd -> Actor: ?M** respectively. Note that **?** acts as a wildcard if it is not followed by any variable (e.g., **Snd -> Actor: ?**) since no specific pattern of the receiving message is expected.
 - It is possible to use different kinds of channels: **->**, ***->**, **->*** and ***->***, which define insecure, authentic, confidential and secure channels, respectively. In ASLan++, there are different types of channel models but we only use the Cryptographic Channel Model (CCM).
 - A section **clauses** defines Horn clauses of the form: **Hcname: head :- body**.
 - In the **goals** section, one can specify LTL goals like **[]predicate(Var)** stating that a particular predicate must always hold over a variable **Var**.
 - The two statements **if(<guard>){positive branch} else {negative branch}** and **select{on:{<guard>}} {positive branch}** are equivalent but **select-on** does not provide the negative branch.
 - **while(<guard>)** loops in ASLan++ are used to define processes waiting for incoming messages (e.g., our database formalization).
 - Assignments are of the form **M:=m** where **M** is a variable and **m** is a constant.
 - There are various types in ASLan++, e.g.: **agent** for entities involved in the communication, **message** for anything that can pass through the network, **text** for atomic messages (i.e., messages that do not contain concatenations, whereas **message** can contain a concatenation of **texts**), and the Boolean type **fact** used for predicates. ASLan++ supports also functions and sets whose types can be anything but **fact**.
 - The main ASLan++ keywords are: **Actor**, **?**, **iknows()**, which is a predicate that represents the intruder knowledge (it stands for *intruder knows*), and

`fresh()`, which is a function that generates a fresh constant value that one can assign to a variable by `M:=fresh()`.

- Finally, we describe here the only type of ASLan++ goal we have used in our examples, i.e., `[](!iknows(M))` where `[]` is the globally LTL (Linear Temporal Logic) operator, `!` expresses a negation, `iknows` is as above and `M` (in our examples) is a constant or the `tuple()` function. Informally, the goal states that, in a safe ASLan++ specification, `M` will never be known by the intruder, otherwise an attack is found.

An ASLan++ specification can be automatically translated (see [3]) into a more low-level ASLan specification, which ultimately defines a transition system $M = \langle \mathbf{S}, \mathbf{I}, \rightarrow \rangle$, where \mathbf{S} is the set of states, $\mathbf{I} \subseteq \mathbf{S}$ is the set of initial states, and $\rightarrow \subseteq \mathbf{S} \times \mathbf{S}$ is the (reflexive) transition relation. The structure of an ASLan specification is composed by six different sections: signature of the predicates, types of variables and constants, initial state, Horn clauses, transition rules of \rightarrow and protocol goals. The content of the sections is intuitively described by their names. In particular, an initial state $I \in \mathbf{I}$ is composed by the concatenation of all the predicates that hold before running any rule (e.g., the agent names and the intruder's own keys). The transition relation \rightarrow is defined as follows. For all $S \in \mathbf{S}$, $S \rightarrow S'$ iff there exist a rule such that

$$PP.NP \& PC \& NC \models [V] \Rightarrow R$$

(where PP and NP are sets of positive and negative predicates, PC and NC conjunctions of positive and negative atomic conditions) and a substitution $\gamma : \{v_1, \dots, v_n\} \rightarrow T_\Sigma$ where v_1, \dots, v_n are the variables that occur in PP and PC such that: (1) $PP\gamma \subseteq [S]^H$, where $[S]^H$ is the closure of S with respect to the set of clauses H , (2) $PC\gamma$ holds, (3) $NP\gamma\gamma' \cap [S]^H = \emptyset$ for all substitutions γ' such that $NP\gamma\gamma'$ is ground, (4) $NC\gamma\gamma'$ holds for all substitutions γ' such that $NC\gamma\gamma'$ is ground and (5) $S' = (S \setminus PP\gamma) \cup R\gamma\gamma''$, where γ'' is any substitution such that for all $v \in V$, $v\gamma''$ does not occur in S .

We now define the translation of the ASLan++ constructs we have considered here. Every ASLan++ entity is translated into a new *state predicate* and added to the section signature. This predicate is parametrized with respect to a *step label* (that uniquely identifies every instance) and it mainly keeps track of the local state of an instance (current values of whose variables) and expresses the control flow of the entity by means of step labels. As an example, if we have the ASLan++ entity

```
entity Snd(Actor, Rcv: agent){
  symbols
  Var: message;
}
```

then the predicate `stateSnd` is added to the section signature and, supposing an instantiation of the entity `new Snd(snd, rcv)`, the new predicate `state_Snd(snd, iid, sl_0, rcv, dummy_message)` is used in transition rules to store all the informations of an entity, where the ID `iid` identifies a particular instance, `sl_0` is the step label, the parameters `Actor`, `Rcv` are replaced with constants `snd`

and `rcv`, respectively, and the message variable `Var` is initially instantiated with `dummy_message`.

Given that an ASLan++ is a hierarchy of entities, when an entity is translated into ASLan, this hierarchy is preserved by a `child(id_1, id_0)` predicate that states `id_0` is the parent entity of `id_1` and both `id_0` and `id_1` are entity IDs.

A variable assignment statement is translated into a transition rule inside the rules section. As an example, if in the body of the entity `Snd` defined above there is an assignment `Var := constant`; where `constant` is of the same type of `Var`, then we obtain the following transition rule:

```
state_Snd(Actor, IID, sl, Rcv, Var)
=>
state_Snd(Actor, IID, succ(sl), Rcv, constant)
```

In the case of assignments to `fresh()`, the variable `Var` is assigned to a new variable.

In the case of a message exchange (sending or receiving statements), the `iknows(message)` predicate is added to the right-hand side of the corresponding ASLan rule. This states that the message `message` has been sent over the network and `iknows` is used because, as is usual, the Dolev-Yao intruder is identified with the network itself.

The last point we describe is the translation of goals focusing only on the LTL goal we have used in our case studies. Goals are translated into attack states containing the negation of the argument of the LTL operator:

```
attack_state authorization :=
    iknows(M)
```

More information on ASLan, ASLan++ and the AVANTSSAR Platform can be found in [4,3].

B.2 ASLan++ skeleton

We now present a general ASLan++ specification that contains all the aspects described in § 4. We start by describing the first part of the skeleton given in Listing 1.10, where we specify agents, variables, constants, facts and uninterpreted functions used in the overall specification (i.e., the entity `Environment`).

Listing 1.10. ASLan++ code of the symbols used in the skeleton of the web app.

```
1 specification SpecificationSkeleton
2 channel_model CCM
3 entity Environment {
4     symbols
5     %entities involved in the communication
6     webapp, database: agent;
7     %DBMS
8     nonpublic inDB(message): fact;
9     nonpublic sanitizedQuery(message): message;
10    nonpublic query(message): message;
11    nonpublic tuple(message): message;
12    nonpublic no_tuple: text;
13    %sql injection payload
14    sqli: text;
```

```

15   nonpublic dashboard: text;
16   nonpublic secureFolder: text;
17   clauses
18   db_hc_ev(M): inDB(M.sqli);

```

Lines 1–3 begin the specification by stating a name (in this case, `SpecificationSkeleton`) and the channel model used (`CCM`), and by introducing the outermost entity (`Environment`). The symbols section of the environment begins in line 4, where we define the constants representing agents involved in the specification (line 6). We only need to represent the web app and the database but no client as we defined in § 4. In lines 8–12, we define predicates, uninterpreted functions and constants as described in § 4.2. They are defined as `nonpublic` to exclude them from the initial knowledge of the DY intruder. In line 14, we define the payload of § 4.1 that the DY intruder can use to perform SQLi. The constant `sqli` is public in order to add it to the DY intruder initial knowledge. In lines 15–16, we define two constants representing two recurring components of web apps: `dashboard`, which represents a user administration page, `secureFolder`, which represents any secure folder (or page) in a web app. In line 17, the keyword `clauses` opens the Horn clauses section in which the Horn clause in line 18 represents `db_hc_ev` of § 4.2.

Our skeleton considers two main subentities, defining the web app (Listing 1.11) and the database (Listing 1.12). These two entities are subentities of `Session` as showed in line 1 of Listing 1.11, but we will come back to the session entity later in this section.

Listing 1.11. ASLan++ code of the web app entity.

```

1   entity Session(Webapp, Database:agent){
2     entity Webapp(Actor, Database:agent){
3       symbols
4       %all the symbols used in the body of this
5       %entity in the body below
6       body{ %write the behavior of the web app.
7         %Every time a message M is sent to the
8         %database do not forget to add a nonce
9         %to avoid spurious replay attacks,
10        %e.g., Actor *->* database: M.nonce;
11        %the same for message received,
12        %e.g., database *->* Actor: ?M.?Nonce }}

```

The `Webapp` entity has an empty `body` as it has no fixed structure and depends on the particular app one is modeling; we add it with comments to help the modeling phase. As in any other ASLan++ entity, the `symbols` section collects the type definition of variables, constants and eventually functions or predicate used in section `body`, where the behavior of the web app must be specified following the formalization in § 4.3.

The `Database` entity follows (almost verbatim) the pseudo-code in Listing 1.1. The while loop in line 6 wraps the entire body content. This defines that the database is listening for incoming communications that match one of the guards of the two select-on in lines 9 and 14. The other two main differences are the usage of select-on instead of an if statement and the introduction of nonces. The first is due to technical reasons: when the ASLan++ specification is translated into a transition system, the semantics of select-on with respect

to an if statement saves one or more transitions (thanks to the absence of the negative branch). The introduction of nonces in any in/out-going message is useful to avoid spurious man-in-the-middle attacks between the web app and the database in accordance with § 4.3.

Listing 1.12. ASLan++ code of the database entity.

```

1  entity Database(WebApp, Actor: agent){
2  symbols
3    NonceWA, NonceDB: text;
4    SQLQuery: message;
5  body{
6    while(true){
7      select{
8        on(WebApp *->* Actor: ?NonceWA.sanitizedQuery(?SQLQuery)):{
9          select{on(SQLQuery = tuple(?)):{
10             NonceDB := fresh();
11             Actor *->* WebApp: no_tuple.NonceDB; } } }
12        on(WebApp *->* Actor:
13           ?NonceWA.query(?SQLQuery)):{
14          select{
15            on(inDB(SQLQuery)):{
16              NonceDB := fresh();
17              Actor *->* WebApp: tuple(SQLQuery).NonceDB; }
18            on(!(inDB(SQLQuery))){
19              NonceDB := fresh();
20              Actor *->* WebApp: no_tuple.NonceDB; }}}}

```

As we mentioned, both the web app and the database are subentities of the session entity. In our skeleton, the only meaning of the session entity is to instantiate its subentities with the proper constants (defined in Listing 1.10) as showed in Listing 1.13.

Listing 1.13. ASLan++ code of the Session body.

```

1  body{ new Webapp(webapp, database);
2        new Database(webapp, database); }

```

Goals (i.e., security properties that we want to check) are defined in Listing 1.14. They are a verbatim copy of the ones described in § 4.4 but we use `iknows` instead of `iknowledge`.

Listing 1.14. ASLan++ code for authentication and data extraction attacks.

```

1  goals
2    authentication: [](!(iknows(dashboard)));
3    data_extraction: [](!(iknows(tuple(?))));

```

At the end of the specification, the `Environment` entity is instantiated as in Listing 1.15.

Listing 1.15. ASLan++ code of the Environment body.

```

body{ new Session(webapp, database); }

```

C ASLan++ Case Studies

In this section, we give the ASLan++ code that implements the behavior of the case studies presented in § 5 long with two case studies coming from WebGoat. The ASLan++ of this section will fill the empty spaces of the ASLan++ skeleton of § B.2.

C.1 Authentication Bypass (WebGoat)

We used Lesson *Stage 1: String SQL injection* from WebGoat which implements a common login scenario for authenticating users through username and password. The web app receives credentials (`Employeeid` and `Password`) along with the IP address of the client (1). A nonce `NonceWA` is generated (2) to ensure a fresh communication (avoiding spurious replay attacks) with the database and then a query with the provided credentials is sent to the database entity (3). The web app waits to receive an answer from the database (5-6) and, if some tuples are received (6), the users list page (`usersList`) is returned to the client along with the `tuple()` function (7). If no tuples are generated (10), then the login page is shown (11).

Listing 1.16. ASLan++ code representing the WebGoat Stage 1 scenario, where `Actor` refers to the web application.

```
1  ? ->* Actor: ?IP.?Employeeid.?Password;
2  NonceWA:=fresh();
3  Actor *->* Database: NonceWA.query(Employeeid.Password);
4
5  select{
6  on(Database *->* Actor: tuple(?SQLquery).?NonceDB):{
7    Actor ->* IP: usersList.tuple(SQLquery);
8  }
9
10 on(DB *->* Actor: no_tuple):{
11   Actor ->* IP : login;
12 }}
```

As goal (Listing 1.17), we check if there exists an execution in which the intruder can obtain access to the `usersList` constant (representing the users list page) without knowing the correct credentials.

Listing 1.17. Authentication goal.

```
[](!(iknows(usersList)));
```

C.2 Data Extraction (WebGoat)

Consider again the WebGoat lesson specified in § C.1. We can use it again to show how to represent a data extraction attack. The behavior of the web application does not change and is the same given in Listing 1.16. As goal (Listing 1.18), we check if there is a way for the intruder to gain knowledge of something which is function on `tuple()`.

The AAT obtained is identical to the one already given in Listing 1.16 for authentication bypass. Not surprising, whenever there is an injection point, the intruder can modify the behavior of the SQL query by using different payloads to obtain different results. In our model, we do not consider concrete payloads and then the attack execution is the same. What allows us to distinguish between the two cases is the goal itself, which gives us information about the intruder's intention.

Listing 1.18. Data extraction goal

```
[](!(iknows(tuple(?))));
```

C.3 Authentication bypass via Data Extraction

The ASlan++ model representing the Joomla! case study in § 5.1 is given in Listing 1.19. A remote user, browsing the web app, visits the Content History component (line 1). The web app generates the nonce `NonceWA` (2) to ensure a fresh communication (avoiding spurious replay attacks) with the database and then sends a query with the user supplied data (3). The web application waits for a response from the database (5). If some tuples are generated (5), the web application sends back to the client the history page along with the `tuple()` function (6). The web application then has two possible ways of authenticating the user (9–27): by using credentials or cookies. If username and password (12) are provided, the web application applies a non-invertible hash function `hash()` to the password, and queries the database to verify the credentials (13). If the credentials are correct (14), the administration panel is sent to the user (15). In case of a cookie session (20), the user provides a cookie that the web application checks querying the database (21). If the cookie is valid (22), the administration panel is sent back to the user (23).

Listing 1.19. ASlan++ code representing the Joomla! scenario, where `Actor` refers to the web application.

```

1  ? ->* Actor: ?IP.(com_contenthistory.history)?.Listselect;
2  NonceWA := fresh();
3  Actor *->* Database: NonceWA.query((com_contenthistory.history).Listselect
   );
4
5  select{on(Database *->* Actor: tuple(?SQLquery)?.NonceDB):{
6    Actor ->* IP: viewHistory.tuple(SQLquery);
7  }}
8
9  % attempt to access a restricted area
10 select{
11   % we are providing correct credentials
12   on(IP ->* Actor: username.?Username.password.?Password):{
13     Actor *->* Database: NonceWA.sanitizedQuery(Username.hash(Password));
14     select{on(Database *->* Actor: no_tuple?.NonceDB ):{
15       Actor ->* IP: adminPanel;
16     }}
17   }
18 }
19 % we are providing a valid session
20 on(IP ->* Actor: cookie.?Cookie):{
21   Actor *->* Database: NonceWA.sanitizedQuery(Cookie);
22   select{on(Database *->* Actor: no_tuple?.NonceDB):{
23     Actor ->* IP: adminPanel;
24   }}
25 }
26 }%end select
27 }%end body
28

```

As goal (Listing 1.20), we check if there exists an execution in which the intruder can access the administration panel represented by the constant `adminPanel`.

Listing 1.20. Authentication bypass for the Joomla! scenario.

```

[ ] (! (iknows(adminPanel)));

```

YAVWA The ASlan++ model representing the YAVWA case study in § 5.1 is given in Listing 1.21. The client sends his personal credentials (**Username** and **Password**) along with his IP address (**IP**) to the web application (1). The web application generates the nonce **NonceWA** (2) to ensure a fresh communication (avoiding spurious replay attacks) with the database and then creates a query that it sends to the database (3). The web application waits for a response of the database with a tuple (4–5). If tuples are generated from the database (5), a dashboard page is returned to the client along with the tuple (6). If no tuples are generated (7), the web application redirects the user to the login page (8). At this point, the web application waits to receive correct credentials (11) for the HTTP basic authentication that will allow the client to access the secure folder **secureFolder** (12). Given that the credentials are the same as the ones stored in the database, we can model them as a function of **tuple()**. We recall, as discussed in § 4.2, that **tuple()** represents an abstraction of any data that can be extracted from the database. This means that whenever a web application requires any data that is in the domain of the database, we can write them as a function of **tuple()**.

Listing 1.21. ASlan++ code representing the YAVWA scenario, where **Actor** refers to the web application.

```

1  ? ->* Actor: ?IP.?Username.?Password;
2  NonceWA := fresh();
3  Actor *->* Database: NonceWA.query(Username.Password);
4  select{
5    on(Database *->* Actor: tuple(?SQLquery).?NonceDB){
6      Actor ->* IP:dashboard.tuple(SQLquery);
7    on(Database *->* Actor: no_tuple.?NonceDB){
8      Actor ->* IP: login;
9    }
10 }
11 select{on(IP ->* Actor: tuple(?)):{
12   Actor ->* IP: secureFolder;
13 }}

```

As a goal (Listing 1.22), we ask if there is an execution in which the intruder reaches **secureFolder**.

Listing 1.22. Authorization goal

```

[](! (iknows(secureFolder)));

```

C.4 Second-Order SQLi (SO)

The ASlan++ model representing the case study in § 5.2 is given in Listing 1.23. The client sends a registration request together with his IP address (**IP**) and his personal credentials (**Username** and **Password**) to the web application (1). The web application generates a nonce **NonceWA** (2) to ensure a fresh communication with the database (to avoid spurious replay attacks) and sends a query containing the client’s credentials along with the nonce to the database (3). The web application waits until it receives a response from the database containing the data resulting from the execution of the query **tuple(?SQLquery)** submitted by the web application (5). The registration process is now completed and the web

application sends back to the client the page `registered` (6). Here, the web application does not forward `tuple()` back to the client because the registration query is an `INSERT` (see § 4.2). The client asks for a page (8) and the web application uses previously submitted values of `Username` and `Password` to execute a new SQL query (9). Here is where the SO takes place; in fact, the variables embedded in the query in (9) will trigger a SO. The database executes the query and sends back the results to the web application (11). Finally, (12), the web application sends to the client the requested page and the `tuple()` (the query submitted by the web application is a `SELECT`).

Listing 1.23. ASLan++ code representing a web application vulnerable to a SO attack where `Actor` refers to the web application.

```

1  ? ->* Actor: registrationRequest.?IP.?Username.?Password;
2  NonceWA := fresh();
3  Actor *->* Database: NonceWA.query(Username.Password);
4
5  select{ on(Database *->* Actor: tuple(?SQLquery).?NonceDB):{
6    Actor ->* IP: registered;
7
8    IP *->* Actor: requestPage;
9    Actor *->* Database: NonceWA.query(Username.Password);
10
11   Database *->* Actor: tuple(?SQLquery).?NonceDB;
12   Actor ->* IP: page.tuple(SQLquery);
13 }}

```

As goal (Listing 1.24), we ask if there is an execution in which the intruder can interact with the web application until he obtains data from the database, i.e.

Listing 1.24. Database data leakage goal.

```
[](! (iknows(tuple(?))));
```

C.5 Concretization phase

We give the the output provided by SQLfast with respect to the Joomla! case study in Listing 1.25.

Listing 1.25. Concretization of the SQLi in the AAT Listing 1.5.

```

1  Just a couple of questions.
2  What's the name of the web app in the ASLan++ specification? [DEFAULT
   webapplication, hit enter twice for default]
3
4  What's the name of the database in the ASLan++ specification? [DEFAULT
   database, hit enter twice for default]
5
6  Can you give me the URI of the web page under test corresponding to:
7  i -> WebApp: i.Username(4).sqli
8  0 - <?> ->* webapplication : com_contenthistory.history.sqli
9  http://target.com/joomla3.4.4/index.php?list[select]=?&view=history&option
   =com_contenthistory
10 -----
11 0 - <?> ->* webapplication : com_contenthistory.history.sqli
12 http://target.com/joomla3.4.4/index.php?list[select]=?&view=history&option
   =com_contenthistory
13
14 Data extraction command

```

```
15  Are there any POST parameters (key=value)? [optional, press enter to skip]
16
17  sqlmap.py -u "https://157.27.244.25/joomla3.4.4/index.php?list[select]=?&
    view=history&option=com_contenthistory" -a
18  -----
```

SQLfast first asks to the modeler the names of the ASLan++ entities representing the web application and the database (1-5). Then, SQLfast automatically detects the SQLi in the AAT and asks for the URL corresponding to the message in the attack trace (lines 6-9). Once provided, SQLfast asks whether there are any POST parameters for performing the request to the give URL (15-16) and finally generates the appropriate command that need to be executed (17).